



INTERRUPTING A MICROPROCESSOR ONLY AFTER A TRANSACTION IS COMPLETE

5

BACKGROUND

In many systems, expansion devices along expansion busses that allow communication with the system processor add capabilities and functions. A common example of these expansion devices is circuits on expansion boards, or cards, that can be inserted into a backplane of a computing system. The expansion bus provides 10 communication between the new card and the system processor. A common example of a bus protocol used in this situation is a Peripheral Component Interconnect (PCI) bus.

The PCI bus is linked to the system bus upon which resides the system memory, the system processor and other system components by a device referred to here as a PCI bridge. In the case of more than one expansion bus, bridges may exist between the expansion busses 15 that are generally referred to as PCI-to-PCI bridges, or P2P bridges. As used here, the term bridge may be used to refer to either type of bridge. Multiple P2P bridges may exist between the expansion card and the system bus.

In most PCI implementations, the PCI card desires to transmit data to the system memory. The PCI card transmits the data desired to the system memory through a bridge. 20 When the PCI card has completed its transmission of the data, it transmits an interrupt along a command path to the processor. The command path is generally much faster than the path taken by the data, referred to here as the data path.

For example, if the data transfer from the card is complete, the data may have been transferred to the bridge between the PCI device and the system memory. The bridge may be 25 handling transactions for several different PCI devices and the data may be temporarily stored in a buffer until it can be transmitted to the system memory. This is just one example of latencies in the data path. The result is that the processor receives the interrupt prior to the data being in memory. The processor may issue a read request for the address of the data and

then stall until all of the data arrives at the address. This reduces the processor efficiency, as the processor essentially becomes idle while waiting on the data.

In one approach, a 'false' or 'dummy' PCI read is issued to the device when the interrupt is received at the microprocessor. When the PCI device receives the read request, it transfers its data to the system memory. When that transfer is complete, the processor is assured of having all of the necessary data in the system memory. While this may alleviate some of the problem, there is still a stall that occurs during the processing of the dummy read cycle.

BRIEF DESCRIPTION OF THE DRAWINGS

10 The invention may be best understood by reading the disclosure with reference to the drawings, wherein:

Figure 1 shows an embodiment of a system having an expansion bus.

15 Figure 2 shows an alternative embodiment of a system having multiple expansion busses.

Figure 3 shows an embodiment of a method to process a transaction received on an expansion bus.

Figure 4 shows a more detailed embodiment of a method to process a transaction received on an expansion bus.

20 Figure 5 shows an alternative, more detailed embodiment of a method to process a transaction on an expansion bus.

DETAILED DESCRIPTION OF THE EMBODIMENTS

Figure 1 shows an embodiment of a system having an expansion bus. The central processing unit 10, or system processor, is connected to a system memory 12, and a bridge device 14 on a local bus. As will be discussed in more detail further, the memory 12 may have a predetermined area reserved for writes from the bridge 14. A memory controller 19 resides at the junction of the local bus and the output of the bridge 14.

One or more expansion busses have devices such as 18 residing on them. As will be discussed further in more detail, there may be several expansion busses and the bridges may bridge between the expansion bus and the local system bus, or the bridges may bridge between two or more expansion busses. There can be several buffers along the path taken by 5 data from the expansion device and the system processor.

The bridges generally operate in a first-in-first-out procedure (FIFO) and may have several devices transmitting data to the system processor. Several devices may have data routing through the bridge, so the devices 'compete' for slots in the processing, or transaction, queue of the processor in the bridge. While waiting to be processed, the data may 10 be stored in a buffer in that bridge. This causes delays in the data path that may result in inefficiencies in system operation.

For example, the device 18 may transmit the data to the system processor along the expansion bus. When the device has released the data, it transmits an interrupt along a different path, referred to here as a command path. The command path is faster than the data 15 path, and the system processor may receive the interrupt before the data is in the system memory. The processor responds to the interrupt as soon as it receives it, but the data may not be in system memory. It has to wait for the data until it is sure that all data is in system memory. While doing so it stalls, meaning that it does no further processing while it waits for the data.

20 An analogy would be a road passing by two checkpoints with checkpoints in between. When a certain car passes the first checkpoint being the bridge, a light or electrical signal is transmitted to the checkpoint at the other end of the road stating that the car has passed. The end checkpoint receives the signal almost immediately. Meanwhile, the car has to stop and be checked at the checkpoints in between, so the end point that received the signal has to wait 25 for the car or cars to reach it. The receiving checkpoint does not know how many cars it is expecting, so it just waits until there are no more.

Current approaches have attempted to address this problem. In one approach, upon receipt of an interrupt at the CPU a ‘dummy’ read transaction is generated by the CPU back to the expansion device that is transmitting the interrupt. When the response to the dummy read returns to the CPU, the CPU can conclude that all the data previously transmitted has 5 arrived. The response to the read request will travel in the data path, and therefore will follow the last portion of the data that was previously transmitted and the transmission of which was what cause the interrupt.

Returning to the analogy above, this would be similar to receiving the signal that the last car has passed the checkpoint and then sending a following car from the receiving 10 checkpoint to the other end and waiting until the following car has returned, thereby verifying that all the original cars have passed by. In this instance, the route traveled includes the bridge checkpoints, such as bridge 14.

The bridge device 14 has a first port 144 that communicates with the expansion bus. A second port 142 communicates with another bus, which may be a system bus or another 15 expansion bus. The bridge processor 140 manages the communications between the two busses. Memory controller 19 resides at the junction of local bus and the bridge output. Memory is attached to the memory controller, which also connects to the local bus and the expansion bus.

An alternative arrangement of a system having expansion busses is shown in Figure 2. 20 The components of this arrangement have been simplified for ease of discussion, but would generally be arranged according to their counterparts in Figure 1. The CPU 10 is connected through a system bus 17 to a first bridge 14. A first expansion bus 17a then connects the first bridge 14 to a second bridge 14b. A second expansion bus 17b then connects the second bridge 14b to the transmitting device 18a. Data being transmitted from the transmitting 25 device 18a must traverse all three busses and both bridges before reaching the system

memory, at which point the CPU could operate upon it. Meanwhile, the interrupt would be transmitted along a command path that is much faster than the data path.

For example, assume device 18a is the one transmitting the data, there may be delays as the data travels from 18a along the expansion bus 17b to bridge 14. The data then has to 5 waits its turn in the queue at 14, before it gets transmitted along the bus 17a to the bridge 14. The data then again has to wait as it is processed through the queue at bridge 14 then transmitted to the system memory along bus 17.

Two devices 18a and 18b may each have an expansion bus, 17b and 17c. This is done to avoid device 18a from interfering with 18b, which may occur if the PCI bus is shared.

10 Problems may occur if a pin is bent or similar issue occurs. Device 14 may also have more expansion busses such as 17d, with more bridges. This may be done to keep expansion bus lengths short.

A change to the process can be implemented by embodiments of the invention. As can be seen in Figure 3, the bridge can be changed to intercept the interrupt from the device 15 18 as shown in 20 of Figure 3. When the bridge intercepts the interrupt, it essentially stops it from reaching the processor. As the bridge intercepting the interrupt is the one just next to the transmitting device, the bridge should have the data prior to the interrupt being received at the bridge. The interrupt is transmitted from the device after the data is sent and there are no latencies in the path between the device and the first bridge.

20 The bridge processes the data transmitted to it as normal, such as inserting it into the transaction queue at 22. The bridge alters its operation at 24 and generates an indicator of completion, where the completion is the completion of the data transmission to the system processor. This indicator is inserted into the queue after the data at 26. This indicator will travel along the data path, rather than the command path.

25 The effect of the transaction being inserted into the queue on the data path is that it will not reach the system processor or memory until the data transmitted from the expansion

device has already been received into system memory. Returning to the previous analogy, when the last car passes the first checkpoint, the first checkpoint sends a vehicle by itself after the last car passes. The receiving checkpoint can go on about its business until the vehicle from the checkpoint arrives.

5 The transaction inserted into the queue may be addressed to a predetermined region of memory, such as 120 shown in Figure 1. In a more detailed view of the process 28 of Figure 3, shown in Figure 4, the transaction is generated at 30. The transaction is inserted into the transaction queue at 32. Eventually, this transaction will result in a write to the predetermined area of memory.

10 When the write occurs, the logic in the memory controller generates an interrupt to the processor. The interrupt is then processed according to one of many ways known in the art. However, the processor now ‘knows’ that the data transmitted from the expansion device is in the memory and can operate on that data as needed. During the period of time between the interrupt being intercepted at the bridge and the write occurring at the system memory, the 15 processor is not stalling, allowing the system to avoid the inefficiencies of a processor stall.

In an alternative approach, the process 28 of Figure 3 may be altered slightly as shown in Figure 5. Instead of the bridge generating a new transaction and inserting it into the queue, the bridge could issue a read request to an address on the expansion device at 40. The bridge would process the response to the read request from the expansion device at 42 after the 20 transmitted data had already been processed, resulting in the same basic effect as the previous alternative. The memory controller logic may have to interrupt the processor with the memory location is updated.

Generally, bridge devices are application specific integrated circuits (ASICs) that are manufactured specifically to act as bridges. In some cases, however, the bridges may be 25 deployed as Field-Programmable Gate Arrays (FPGAs), general-purpose processors, or digital signal processors. In any of these implementations, it is possible that embodiments of

the invention may take the form of an article of machine-readable code, where the machine is the bridge device. The code may be loaded into the bridge device, thereby upgrading the instructions to be executed by the bridge device. The instructions, when executed, cause the machine to perform the methods of the invention.

5 However implemented, the embodiments of the invention allow the processor to continue to perform tasks without stalling to process reads of data that has not yet arrived into memory. This increases the efficiency of the system, as well as allowing more certainty as to the completeness of data being transmitted. The transaction that is inserted after the data not only ensures that the processor will pause to read data already present, rather than data not yet 10 arrived, but that the data received is complete. Complete as used here means that any data that is expected has arrived, since the transaction would arrive after any data transmitted.

Thus, although there has been described to this point a particular embodiment for a method and apparatus for processing a transaction from an expansion device, it is not intended that such specific references be considered as limitations upon the scope of this 15 invention except in-so-far as set forth in the following claims.